

What is claimed is:

1. A method for safeguarding data in a multi-cluster storage server, comprising:
  - maintaining respective data structures that describe data stored in respective caches of the clusters;
  - wherein for each cluster, the data is provided in a plurality of track images of the respective cache, and the respective data structures include track control blocks that describe the data in the track images; and
  - following a failure of one of the clusters, using the track control blocks of the surviving cluster to rebuild the data structures for the respective cache.
2. The method of claim 1, further comprising:
  - verifying an integrity of the track control blocks of the surviving cluster before using them to rebuild the data structures for the respective cache.
3. The method of claim 2, wherein:
  - the integrity of the track control blocks of the surviving cluster is verified following a reboot in the surviving cluster that follows the failure.
4. The method of claim 2, further comprising:
  - maintaining a non-volatile memory with information identifying track images containing modified data, and maintaining information in the track control blocks indicating whether the track images contain modified data;
  - wherein verifying the integrity of the track control blocks comprises determining whether the information maintained in the non-volatile memory corresponds with the information maintained in the track control blocks.

5. The method of claim 4, further comprising:  
updating the information maintained in the non-volatile memory and in the track control blocks when data is added to, or removed from, the respective cache.
6. The method of claim 2, further comprising:  
placing on a free list the track images described by track control blocks whose integrity cannot be verified to enable reuse.
7. The method of claim 1, wherein:  
the data structures for the respective cache that are rebuilt include a least recently used (LRU)/most recently used (MRU) list indicating a probability of a read access by a host of the data provided in the plurality of track images of the respective cache.
8. The method of claim 7, wherein:  
the LRU/MRU list is rebuilt by adding track control blocks that are determined to be valid and describe track images having modified data to the LRU/MRU list.
9. The method of claim 7, wherein:  
the track control blocks include a backward pointer identifying a previous track control block in the LRU/MRU list, and a forward pointer identifying a next track control block in the LRU/MRU list.
10. The method of claim 1, wherein:  
the data structures for the respective cache that are rebuilt include a scatter table that outputs information identifying locations of the track images in the respective cache.

11. The method of claim 10, wherein:  
the scatter table outputs the information identifying locations of the track images in the respective cache based on hashed values of identifiers of the track control blocks.
12. The method of claim 10, wherein:  
the scatter table is rebuilt by adding track control blocks that are determined to be valid and describe track images having modified data to the scatter table.
13. The method of claim 1, wherein:  
rebuilding the data structures for the respective cache enables the data stored therein to be accessed.
14. A computer system for safeguarding data in a multi-cluster storage server, comprising:  
memory in which respective data structures that describe data stored in respective caches of the clusters is maintained;  
wherein for each cluster, the data is provided in a plurality of track images of the respective cache, and the respective data structures include track control blocks that describe the data in the track images; and  
a processor for using the track control blocks of the surviving cluster to rebuild the data structures for the respective cache, following a failure of one of the clusters.
15. The computer system of claim 14, wherein:  
the processor verifies an integrity of the track control blocks of the surviving cluster before using them to rebuild the data structures for the respective cache.
16. The computer system of claim 14, further comprising:

a non-volatile memory maintained with information identifying track images containing modified data, and maintaining information in the track control blocks indicating whether the track images contain modified data;

wherein the processor verifies the integrity of the track control blocks by determining whether the information maintained in the non-volatile memory corresponds with the information maintained in the track control blocks.

17. The computer system of claim 14, wherein:

the data structures for the respective cache that are rebuilt include a least recently used (LRU)/most recently used (MRU) list indicating a probability of a read access by a host of the data provided in the plurality of track images of the respective cache.

18. The computer system of claim 14, wherein:

the data structures for the respective cache that are rebuilt include a scatter table that outputs information identifying locations of the track images in the respective cache.

19. A computer program product for safeguarding data in a multi-cluster storage server, comprising:

computer code devices configured to cause a computer to: (a) maintain respective data structures that describe data stored in respective caches of the clusters;

wherein for each cluster, the data is provided in a plurality of track images of the respective cache, and the respective data structures include track control blocks that describe the data in the track images; and

(b) following a failure of one of the clusters, use the track control blocks of the surviving cluster to rebuild the data structures for the respective cache.

20. The computer program product of claim 19, wherein:

the computer code devices are configured to cause a computer to verify an integrity of the track control blocks of the surviving cluster before using them to rebuild the data structures for the respective cache.

21. The computer program product of claim 20, wherein:

the computer code devices are configured to cause a computer to maintain a non-volatile memory with information identifying track images containing modified data, and maintain information in the track control blocks indicating whether the track images contain modified data;

wherein verifying the integrity of the track control blocks comprises determining whether the information maintained in the non-volatile memory corresponds with the information maintained in the track control blocks.

22. The computer program product of claim 19, wherein:

the data structures for the respective cache that are rebuilt include a least recently used (LRU)/most recently used (MRU) list indicating a probability of a read access by a host of the data provided in the plurality of track images of the respective cache.

23. The computer program product of claim 19, wherein:

the data structures for the respective cache that are rebuilt include a scatter table that outputs information identifying locations of the track images in the respective cache.

24. A method for safeguarding data in a multi-cluster storage server, comprising:

maintaining respective data structures that describe data stored in respective caches of the clusters;

wherein for each cluster, the data is provided in a plurality of track images of the respective cache, and the respective data structures include track control blocks that describe the data in the track images;

following a failure of one of the clusters, using the track control blocks of the surviving cluster to rebuild the data structures for the respective cache; wherein:

the data structures for the respective cache that are rebuilt include:

(a) a least recently used (LRU)/most recently used (MRU) list indicating a probability of a read access by a host of the data provided in the plurality of track images of the respective cache; and

(b) a scatter table that outputs information identifying locations of the track images in the respective cache.

25. A computer system for safeguarding data in a multi-cluster storage server, comprising:

memory in which respective data structures that describe data stored in respective caches of the clusters is maintained;

wherein for each cluster, the data is provided in a plurality of track images of the respective cache, and the respective data structures include track control blocks that describe the data in the track images; and

a processor for using the track control blocks of the surviving cluster to rebuild the data structures for the respective cache, following a failure of one of the clusters; wherein:

the data structures for the respective cache that are rebuilt include: (a) a least recently used (LRU)/most recently used (MRU) list indicating a probability of a read access by a host of the data provided in the plurality of track images of the respective cache; and

(b) a scatter table that outputs information identifying locations of the track images in the respective cache.

26. A computer program product for safeguarding data in a multi-cluster storage server, comprising:

computer code devices configured to cause a computer to: (a) maintain respective data structures that describe data stored in respective caches of the clusters;

wherein for each cluster, the data is provided in a plurality of track images of the respective cache, and the respective data structures include track control blocks that describe the data in the track images; and

(b) following a failure of one of the clusters, use the track control blocks of the surviving cluster to rebuild the data structures for the respective cache; wherein:

the data structures for the respective cache that are rebuilt include:

(i) a least recently used (LRU)/most recently used (MRU) list indicating a probability of a read access by a host of the data provided in the plurality of track images of the respective cache; and

(ii) a scatter table that outputs information identifying locations of the track images in the respective cache.

27. A method for safeguarding data in a multi-cluster storage server, comprising:

maintaining respective data structures that describe data stored in respective caches of the clusters;

wherein for each cluster, the data is provided in a plurality of track images of the respective cache, and the respective data structures include track control blocks that describe the data in the track images;

maintaining a non-volatile memory with information identifying track images containing modified data, and maintaining information in the track control blocks indicating whether the track images contain modified data;

following a failure of one of the clusters, using the track control blocks of the surviving cluster to rebuild the data structures for the respective cache; and

verifying an integrity of the track control blocks of the surviving cluster before using them to rebuild the data structures for the respective cache by determining whether the information maintained in the non-volatile memory corresponds with the information maintained in the track control blocks.

28. A computer system for safeguarding data in a multi-cluster storage server, comprising:

memory in which respective data structures that describe data stored in respective caches of the clusters is maintained;

wherein for each cluster, the data is provided in a plurality of track images of the respective cache, and the respective data structures include track control blocks that describe the data in the track images;

a non-volatile memory maintained with information identifying track images containing modified data, and maintaining information in the track control blocks indicating whether the track images contain modified data; and

a processor for using the track control blocks of the surviving cluster to rebuild the data structures for the respective cache, following a failure of one of the clusters;

wherein the processor verifies an integrity of the track control blocks of the surviving cluster before using them to rebuild the data structures for the respective cache by determining whether the information maintained in the non-volatile memory corresponds with the information maintained in the track control blocks.



29. A computer program product for safeguarding data in a multi-cluster storage server, comprising:

computer code devices configured to cause a computer to: (a) maintain respective data structures that describe data stored in respective caches of the clusters;

wherein for each cluster, the data is provided in a plurality of track images of the respective cache, and the respective data structures include track control blocks that describe the data in the track images;

(b) maintain a non-volatile memory with information identifying track images containing modified data, and maintain information in the track control blocks indicating whether the track images contain modified data;

(c) following a failure of one of the clusters, use the track control blocks of the surviving cluster to rebuild the data structures for the respective cache; and

(d) verify an integrity of the track control blocks of the surviving cluster before using them to rebuild the data structures for the respective cache by determining whether the information maintained in the non-volatile memory corresponds with the information maintained in the track control blocks.